

EBOOK

Pepperdata Real-Time Cost Optimization for Data-Intensive Workloads on Amazon EMR and EKS



aws
PARTNER

Pepperdata: Your Trusted Partner

Pepperdata is the only cost optimization solution that delivers significant cost savings continuously and in real time on Amazon EMR and Amazon EKS with no application changes or manual tuning. We manage some of the largest and most complex clusters in the world, at top enterprises such as Citibank, IQVIA, Autodesk, Securonix, Royal Bank of Canada, and companies in the Fortune 5.

Driving Savings for Global Enterprises



The Cost Savings Imperative

As evidenced by the expected growth of the FinOps market from an estimated \$800 million in 2023 to \$2.8 billion by 2028 (a CAGR of 18.8 percent), cloud cost optimization has become a clear business imperative.

One third of enterprises expect to exceed their cloud budget by 40 percent or more, which is a difficulty in any market climate, but especially in contracting markets. As a result, customers are increasingly seeking to reduce costs in cloud environments such as Amazon EMR and Amazon EKS.

Pepperdata's Secret Sauce: How Pepperdata Reduces Costs

With over a decade of experience and patented algorithms, Pepperdata Capacity Optimizer addresses the market need to wrangle cloud costs and helps make the cloud an even more viable and attractive option for large and small enterprises alike.

Capacity Optimizer uses Machine Learning (ML) to provide [Continuous Intelligent Tuning](#) that works in real time in the background to safely eliminate application waste and

reduce costs. **No other solution delivers cloud cost optimization immediately, autonomously, continuously, and in real time with no manual tuning or application code changes.**

The time and money Pepperdata saves over manually tweaking applications can be banked, repurposed for new product innovation, and/or redirected to higher value cloud services.

Pepperdata installs in under 30 minutes in most enterprise environments. We guarantee a minimum 100% ROI, with a typical ROI between 100% and 660%.



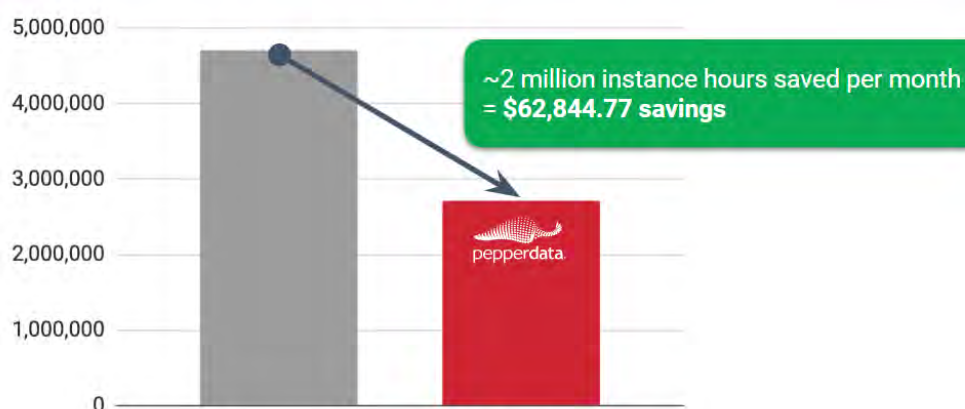
Autonomous Savings in Real Time

Pepperdata Capacity Optimizer software is the “easy button” for cloud cost optimization for Amazon EMR and Amazon EKS.

Capacity Optimizer goes to work immediately once installed. Within an hour, you can automatically visualize measurable cost savings in the Pepperdata-provided customer dashboard which

translates directly to reduced cloud bills at the end of the month. **Capacity Optimizer** does not need to be manually trained, tuned, or retrained on workload patterns. Instead, **it uses ML to constantly calibrate your ever-changing environment in real time to keep your clusters running at their sweet spot of peak utilization**, reducing both instance hours and cost.

42.5% Reduction in Normalized Instance Hours for Amazon EMR on EKS

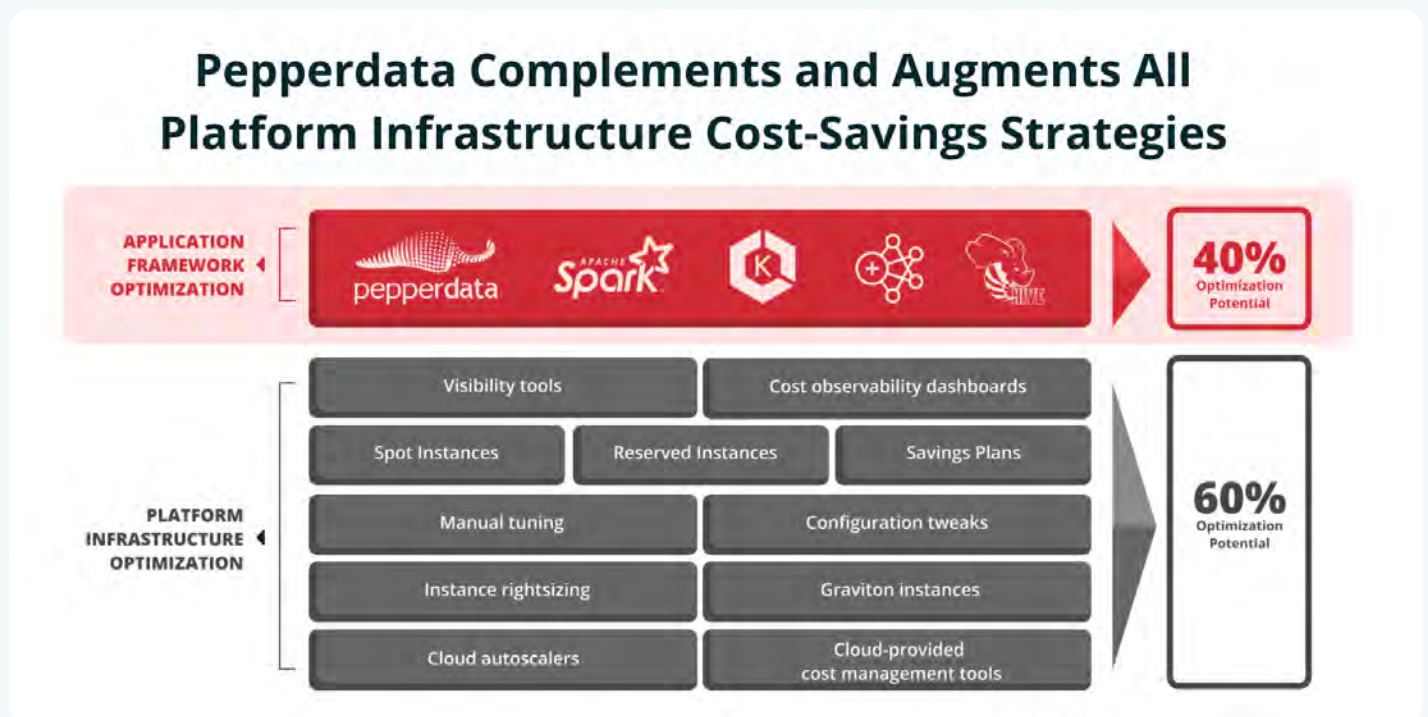


* Actual customer data from October 2023

Application-Level Savings

Most customers today optimize their environments with a combination of Platform Level cloud cost optimization strategies such as Savings Plans, Reserved Instances, Instance Rightsizing, and ongoing manual tuning of applications (see graphic below).

Pepperdata Capacity Optimizer works at a different level. Capacity Optimizer operates at the Application Framework Level to minimize the waste inherent in applications like Apache Spark and thus complements and augments these Platform Level cloud cost optimization strategies.




Maximize the Value of Apache Spark Workloads

Pepperdata Capacity Optimizer autonomously and continuously [reclaims waste from Spark applications](#) in real time, empowering you to:



Reduce Costs:

Save an average of 30-47 percent on Spark workload costs on Amazon EMR and Amazon EKS.



Optimize Spark Clusters for Efficiency:

Minimize (or eliminate) waste in Spark to run more applications without additional spend.




Eliminate Manual Tweaking and Tuning:

Free developers from the tedium of managing individual apps so they can focus on innovation.

Capacity Optimizer is able to do this by:



Automatically identifying in real time where more Spark applications can be run.



Enabling the scheduler to more fully utilize available resources before adding new nodes or pods.



Using Machine Learning (ML) to constantly calibrate your ever-changing environment in real time to keep your clusters running at the sweet spot of peak utilization.

The result: **Spark CPU, memory, and I/O resources are automatically optimized to reduce costs and increase utilization**, enabling more apps to be launched at cost savings of up to 47 percent.

The Pepperdata Dashboard shows actual Realized Savings and Potential Savings at the cluster level so teams know exactly how much they are saving. You can tailor optimization levels by cluster and then view savings in real time.



As one proof point, **Capacity Optimizer empowered Autodesk to boost the performance of their huge Spark jobs and reduce their Amazon EMR costs by 50 percent.**

Pepperdata Customers Enjoy Cost Savings Every Day

Daily Cost Savings from Using Pepperdata

FORTUNE 10 ENTERPRISE	\$13,732
LARGE SECURITY INTELLIGENCE SV	\$7,191
MULTINATIONAL SOFTWARE VENDOR	\$5,285

Reducing Cost in Some of the World's Most Complex and Highly Scaled Clusters

Pepperdata's customers include top enterprises running some of the largest, most complex, and highly-scaled clusters in the world such as Citibank, IQVIA, Autodesk, Securonix, Royal Bank of Canada, and those in the Fortune 5.

The following case studies highlight benefits Pepperdata delivers:

- **Autodesk Reduces Amazon EMR Costs by 50% and Boosts Performance with Pepperdata** ([public reference](#))
- **Extole Exceeded Expectations for Cloud Cost and Resource Optimization with Pepperdata** ([public reference](#))
- **Global Technology Firm Realizes up to 25% Cost Savings on Amazon EMR with Pepperdata** ([anonymized public reference](#))

Detailed Observability Into Spark Applications

Pepperdata's legacy of delivering Spark observability dashboards means that customers using Pepperdata Capacity Optimizer also automatically enjoy single-pane-of-glass visibility into the detailed aspects of Spark application performance.

Pepperdata's dashboards provide advanced visibility into troubleshooting Spark applications and clusters, including Spark metrics, containers, clusters, pods, nodes, and workflows. In addition, Pepperdata reports on detailed usage attribution for chargeback/showback.

This information can be vital to controlling costs [as enterprises migrate their Spark workloads to the cloud](#).

Supported Technologies

- Amazon EMR for EC2
- Amazon EMR for EKS
- Amazon EKS
- Apache Spark
- Kubernetes

Get Started Saving with Pepperdata

Pepperdata installs in under an hour in most environments via a bootstrap script (Amazon EMR) or a Helm chart (Amazon EKS) and immediately begins to extract up to 47 percent in additional resource capacity from even the most highly tuned clusters. **In fact, Pepperdata works best in complicated clusters with many contending applications, long queues, and variable workload characteristics and datasets sizes.**

With Capacity Optimizer, you will immediately begin to realize:

- **Reduced cost and inefficiencies in applications**
- **Optimal price/performance cost outcomes**
- **Overall optimization of cloud infrastructure spend, allowing customers to redirect spend to higher value cloud services**

You can get started with a **[FREE waste assessment](#)** or a no-cost trial **[through the AWS Marketplace](#)**. For a **customized demo**, please contact us at sales@pepperdata.com.

Special Offer

Get started today with a free, 2-day savings assessment of your Amazon EMR or EKS environment. Visit pepperdata.com or contact us at sales@pepperdata.com.

